

Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe® Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

DISCLAIMER: THE MATERIALS ARE PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND INCLUDING WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT OF INTELLECTUAL PROPERTY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT SHALL INTEL OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION, LOSS OF INFORMATION) ARISING OUT OF THE USE OF OR INABILITY TO USE THE MATERIALS, EVEN IF INTEL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME JURISDICTIONS PROHIBIT THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT APPLY TO YOU. INTEL FURTHER DOES NOT WARRANT THE ACCURACY OR COMPLETENESS OF THE INFORMATION, TEXT, GRAPHICS, LINKS OR OTHER ITEMS CONTAINED WITHIN THESE MATERIALS. INTEL MAY MAKE CHANGES TO THESE MATERIALS, OR TO THE PRODUCTS DESCRIBED THEREIN, AT ANY TIME WITHOUT NOTICE. INTEL MAKES NO COMMITMENT TO UPDATE THE MATERIALS.

Table of Contents

(Click on page number to jump to articles)

COVER STORY

An Introduction to Peer-to-Peer Computing	3
---	---

COLUMN

From the Editor.....	8
----------------------	---

DEPARTMENTS

DESKTOP

Intel® D810E2CB Board Makes a True Value Platform	9
---	---

Take Your System's Temperature with TTB	12
---	----

INITIATIVES AND TECHNOLOGIES

CNR Audio Solutions Reduce Costs for OEMs	15
---	----

IDF Spring Conference Takes Broad Industry View	19
---	----

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

An Introduction to Peer-to-Peer Computing

David Barkai
Peer-to-Peer Architecture Group
Microcomputer Research Lab
Intel Corporation

Overview

There is a revolution underway that represents a new computing model for the Internet. Its impact will be comparable to the introduction of the graphical-interface browser.

This revolution is being sparked by a phenomenon known as peer-to-peer (P2P) computing, a usage model that has captured the popular imagination with the highly publicized success of Napster*, Scour*, the Freenet Project, Gnutella*, and SETI@home*.

P2P computing can be simply defined as the sharing of computer resources and services by direct exchange. It's an idea that has achieved considerable traction with mainstream computer users and members of the PC industry. Here are three examples:

- The Napster MP3 music file sharing application went live in September 1999, and attracted more than 20 million users by mid-2000.
- By early 2001, the SETI@home program, which uses distributed processing to analyze radio telescope data, had attracted more than 2.6 million users who had donated over 500,000 years of CPU time to the hunt for extraterrestrial intelligence.
- In October 2000, more than 350 individuals representing about 250 companies and organizations gathered for the first meeting of the Peer-to-Peer Working Group.

P2P computing is enabling a new wave of applications with the potential to revolutionize the way computers are used. Developers should keep this new computing model in mind when creating new applications. Some ways to keep up with the latest technical developments include participating in the Peer-to-Peer Working Group, and attending the Peer-to-Peer panels and tracks at the Intel Developer Forum Conference (IDF) Spring 2001, as well as other conferences dedicated to P2P.

More than Client/Server

P2P computing provides an alternative to the traditional client/server architecture. While employing the existing network, servers, and clients infrastructure, P2P offers a computing model that is orthogonal to the client/server model. The two models coexist, intersect, and complement each other.

In a client/server model, the client makes requests of the server with which it is networked. The server, typically an unattended system, responds to the requests and acts on them.

With P2P computing, each participating computer, referred to as peer, functions as a client with a layer of server functionality. This allows the peer to act both as a client and as a server within the context of a given application. P2P applications build on such functions as storage, computations, messaging, security, and file distribution, when handled through direct exchanges between peers.

A peer can initiate requests, and it can respond to requests from other peers in the network. The ability to make direct exchanges with other users liberates P2P users from the traditional dependence on central servers. Users have a higher degree of autonomy and control over the services they utilize.

One of the greatest benefits of P2P computing is community. P2P makes it possible for users to organize themselves into ad hoc groups that can efficiently and securely fulfill requests, share resources, collaborate, and communicate. As P2P computing evolves, we can anticipate the emergence of a wide variety of these online communities.

Figure 1 illustrates the difference between the client/server and the P2P models. It also illustrates how the two models can be superimposed and coexist with each other. In the client/server model every exchange and communication goes through, and is managed by, a central server. In the P2P computing model, the peer systems communicate and exchange directly. Some P2P applications may, at times, also use servers. The overall effect of P2P computing is to take computing to the edges of the net.

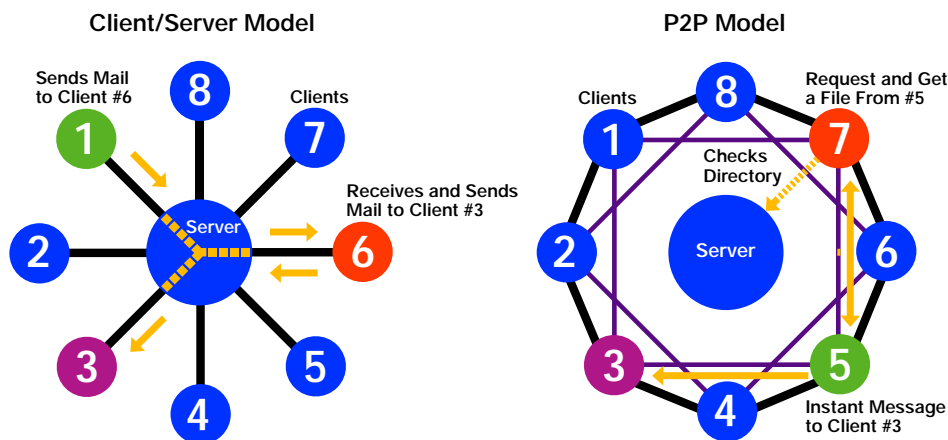


Figure 1. Client Server and P2P Models

What is P2P?

While a peer is a computer that behaves as a client in the client/server model, it also contains an additional layer of software that allows it to perform server functions. The peer computer can respond to requests from other peers. The scope of the requests and responses, and how they are executed, are application-specific.

Typically, there will be a request for access to resources that belong to the other peer. The request may be for information on content and files, or for a file to be read or copied, computations to be performed, or a message file to be passed on to others.

When "computing" in P2P computing is used as a noun, it refers to a framework or computing model. This framework provides the capabilities that allow peers to directly interact. An important characteristic of the direct-interaction capability is that the computing environment becomes decentralized.

When the "computing" in P2P computing is used as a verb, it refers to what we do with P2P framework. Many end-user applications become possible through the P2P services, including storage, computations, messaging, security, distribution, and more. What unifies these application types is sharing of resources with some form of collaboration.

Some P2P evangelists draw a distinction between so-called "pure P2P computing" and "hybrid P2P." The term "pure P2P computing" refers to a model, such as Freenet, where all participating computers are peers. No central server is used to control, coordinate, or manage the exchanges among the peers.

In the "hybrid P2P" computing model, the application relies on a central server to perform some of the required functions. The degree of involvement varies with the application. For example, Napster requires the user to first connect to a control server, where the directory of all available files is stored.

Why P2P?

The P2P computing model offers a number of compelling advantages to individual users and large organizations. These advantages can be divided into technical benefits and social appeal.

Technically, P2P provides the opportunity to make use of vast untapped resources that go unused without it. These resources include processing power for large-scale computations and enormous storage potential.

P2P allows the elimination of the single-source bottleneck. P2P can be used to distribute data and control and load-balance requests across the Net. In addition to helping optimize performance, the P2P mechanism also may be used to eliminate the risk of a single point of failure. When P2P is used within the enterprise, it may be able to replace some costly data center functions with distributed services between clients. Storage, for data retrieval and backup, can be placed on clients. In addition, the P2P infrastructure allows direct access and shared space, and this can enable remote maintenance capability.

Much of the wide appeal of P2P is due to social and psychological factors. For example, users can easily form their own autonomous online communities at the edge of the Net, and run them as they collectively choose. Many of these P2P communities will be ever changing and dynamic in that users can come and go, or be active or not. Other users will enjoy the ability to bypass centralized control. P2P makes users autonomous.

Why Now?

The concept that underlies the P2P computing model, that is, direct exchange and sharing of files and cycles between nodes, is not a new computing concept. So it is reasonable to ask why the P2P model has burst on the scene at this time.

Several essential developments have made P2P for the masses possible. The ubiquity of connected computers has come close to enabling anywhere, anytime access to the Net and its resources. In addition, there is a critical mass of computer users. Improvements in communications bandwidth, still on a fast-track growth curve, make it possible to move large amounts of data and rich media content from one location to another. And today's PCs are sufficiently robust, in terms of processing power and storage capacity, to handle the extra services required in a P2P environment.

Another factor that has accelerated adoption is the emergence of complementary technologies, including recent advances in wireless and software agents that provide more avenues for interesting P2P applications.

While all of these are necessary conditions for P2P computing, something more is required. History shows that a trigger is needed for a new technology to really take off. In the computer industry this is often referred to as a "killer app." The electronic spreadsheet triggered the proliferation of the PC, and the Mosaic* browser triggered the transformation of the Internet into the World Wide Web.

The P2P computing model found its trigger with Napster, followed by Gnutella. Their huge popularity got everyone talking about P2P and helped further stimulate other P2P applications such as Freenet and SETI@home. By the fourth quarter of 2000, over 100 companies and numerous research projects were engaged in P2P computing.

What It's Good For

What are some of the areas where P2P application can thrive? Here are a few examples:

- *Community Web network.* Any group with specific common interests, including a family or hobbyists, can use lists and a Web site to create their own intranet.
- *e-Business.* P2P can add new capabilities, including connecting and enabling the links of a supply chain, distributing information, content, or software more effectively, and keeping information items on their original node with a central directory or a search capability.
- *Gaming.* A P2P infrastructure provides a natural foundation for the development of online community games that are not centrally controlled. Developers can focus on game features instead of the interface to the communications protocol.
- *Search engines.* Fresh, up-to-date information can be found by searching directly across the space where the desired item is likely to reside.
- *Virus protection.* Relationships among the nodes on the P2P community allow for collaboration in virus detection and warning, as well as automatic quarantining of the community against further attacks.
- *Edge services.* There are instances when it is desirable to place the data, prior to its use, closer to the client requesting it. Online training modules that contain video segments, for example, provide the desired effect when the large data files are located close to the online trainee. Multiple clients offering storage space can provide more flexible and reliable service compared to a server.
- *Collaborative development.* The scope can range from developing software products to composing a document to applications like rendering graphics.

What's Needed

While the potential of P2P computing has been demonstrated, and the interest level of the industry is extremely high, several issues remain to be resolved. To date, P2P applications cannot easily communicate with other P2P applications, and each has its own set of basic services and plug-ins. Applications that originate from different sources are usually not integrated and cannot communicate with one another. Even more significant is the fact that developers of new applications cannot readily benefit from the experience of other P2P developers.

As a result, application developers spend much of their time and energy developing services that have already been successfully resolved by others. This is time and energy that could be better spent on adding features and functionality to new applications.

This issue will be resolved when the industry collaborates on solutions that enable interoperability between P2P applications. This interoperability can be achieved through a common set of services that provide the functionality needed for P2P computing above the functionality provided by the peer's operating system. These common services can be thought of as a middleware layer. One of the principal advantages of a common middleware is that application developers will no longer be required to keep creating the same basic services over and over again. Interoperability means that P2P applications can communicate across different software environments. Peers running Windows* or Linux* or any other operating system can share P2P applications. Applications using different programming languages can communicate and be integrated via a common middleware. And common middleware provides a mechanism and an infrastructure for incorporating devices other than PCs and servers. Such devices include wireless and handheld products and various network appliances.

Figure 2 shows a P2P middleware layer that fits between a peer's local operating system and application interfaces that the P2P application can directly call. The creation of standards and a common infrastructure has benefits for users and developers alike.




Figure 2. The P2P Stack

Looking Forward

The P2P computing revolution is happening, and it's big. The success of pioneering P2P applications proves that decentralized computing has enormous appeal for millions of users and obvious benefits for application developers. In addition to "thinking peer-to-peer" during the development process, it is important for developers to understand the power of P2P technology. One of the best ways to do this is to get involved with the Peer-to-Peer Working Group.

The Peer-to-Peer Working Group is an industry consortium dedicated to the advancement of infrastructure standards for P2P computing. The goal of the group is to develop infrastructure standards for interoperability, security, performance, management, and privacy that will enable P2P computing to become ubiquitous. The working group will determine areas for standardization, work to develop specifications, and promote adoption of these specifications as standards throughout the computer industry. Participation is open to anyone interested in peer-to-peer computing.

Developers should also stay tuned for the latest information from Intel, made available through the Intel Developer Forum (IDF) Conference and in future issues of *Intel Developer Update*.

Summary/Survey

For 20 years people in academia and research have experimented with peer-to-peer computing technologies, mostly in relatively small and controlled environments. Recently we have seen P2P applications that have captured the imagination of millions of users.

The idea of sharing and collaborating via the Internet extends to dividing the load of performing large computations, collaborating in creating media or software, conversing online and directly, and organizing into online communities. There are many more intriguing possibilities, limited only by the imagination of application developers and users. P2P computing is evolving into a phenomenon whose impact will be no less than that which we experienced when the graphical-interface browser and when the client/server computing model were introduced.

The physical infrastructure of computers and Internet communications is in place, and the concept of P2P in the mainstream of computing has been demonstrated. The time has come to move from the mode of early experimentation to that of a solid computing model.

More Info

Membership in the Peer-to-Peer Working Group is strongly suggested for any developer interested in creating applications for this revolutionary Internet computing model. See the Peer-to-Peer Working Group home page for membership information and links to information and industry events.

By attending the Intel Developer Forum Conference Spring 2001, you can hear industry leaders share their vision of P2P standards-setting and technical challenges. IDF technical tracks will provide a detailed look at a variety of important issues including security, usage and application examples, and technical challenges.

Future *Intel Developer Update* articles will explore the architecture and infrastructure of P2P computing, and examine some of the technologies and services needed by P2P developers, including security, privacy, resource management, resource sharing, and the advancement of interoperability.

Author Bio

David Barkai joined Intel in 1996. Prior responsibilities at Intel include development of IDF Conference content and work in the Microcomputer Software Lab, where he focused on applications for Intel® Architecture workstations. In addition, David has worked in the area of scientific and engineering supercomputing applications and conducted pioneering work on vector processors. He holds a Ph.D. in theoretical physics and has published more than 20 papers in the areas of physics, numerical methods, and computer applications and architectures.

Column

From the Editor

Donna Loveland
Managing Editor
Intel Developer Update Magazine
Intel Corporation

Column

The Intel Developer Forum's Spring Conference is coming at the end of the month—February 28 through March 1, to be exact—and this issue gives you a look at what's ahead. Not only do we have an article detailing the Conference itself, we also have a selection of articles on just a few of the topics that will be explored in more than 250 IDF labs and sessions.

Be sure to take our opinion poll on peer-to-peer computing, this month's cover story. The poll is a new feature we're launching with this issue, and we're looking forward to seeing what you think.

An Introduction to Peer-to-Peer Computing—Industry interest is building behind peer-to-peer computing, a major new usage model for the Internet. The opportunity is for new types of applications; the challenge for the industry is achieving interoperability.

IDF Spring Conference Takes Broad Industry View—The Intel Developer Forum (IDF) Conference, Spring 2001 offers in-depth technical training, panel discussions, case studies, demos, hands-on labs, and technology overviews for a wide range of hardware and software aspects of client, server, and communication technologies.

Take Your System's Temperature with TTB—Intel's ATX and MicroATX thermal test boards, once proprietary, are now available to system and chassis designers and heat sink developers. The boards enable thermal solution evaluation without instrumentation or data acquisition systems.

Intel® D810E2CB Desktop Board Makes a True Value Platform—The new Intel® D810E2CB Desktop Board is a small, scalable platform with all the essentials developers need to create innovative PC designs and embedded applications.

CNR Audio Solutions Reduce Costs for OEMs—Communication and Networking Riser solutions use the motherboard's AC '97 interface along with processor power and CPU-based DSP software to provide a solution equivalent to PCI add-in cards at or below the PCI add-in card price point.

Next month, be sure to visit *Intel Developer Update* for in-depth information from the IDF Spring Conference.

Enjoy.

Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Platform Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held technical and marketing positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in Expository Writing from the University of Iowa.

Departments

Desktop

Intel® D810E2CB Board Makes a True Value Platform

Sandeep Brahmarouthu
Product Marketing Engineer
DPG Boards/Systems Marketing
Intel Corporation

Overview

A wide range of customers, from home users to IT managers, are searching for new systems that can deliver the performance they require, while lowering costs and saving space. The new Intel® D810E2CB Desktop Board is a small, scalable platform with all the essentials developers need to create innovative PC designs and embedded applications.

What Makes a True Value Platform?

A true value platform has to provide system integrators with the design flexibility and system efficiency they need to create cost-effective designs. It has to incorporate leading technology, feature rock-solid quality, and provide fast integration for time-to-market success.

Based on the new Intel® 810E2 chipset and the Flex ATX form factor, the D810E2CB Desktop Board (See Figure 1) is designed to provide this flexibility. It supports Intel® Pentium® III processors with 100/133-MHz system bus and Intel® Celeron™ processors with 66/100-MHz system bus to deliver efficient performance and scalability. It features the second-generation I/O Controller Hub (ICH2) with ATA 100 hard disk support, four USB ports, and an optional network solution to provide the capabilities users require. With Crystal* sound, two PCI slots, and Intel® graphics already on board, the Desktop Board is a cost-effective platform that addresses both value and mainstream markets.




Figure 1. The Intel Desktop Board D810E2CB is a true value platform that measures just 9.0 x 7.5 inches.

Design Flexibility and System Efficiency

The D810E2CB Desktop Board supports system implementations ranging from sealed, highly integrated solutions to fully featured systems with up to two PCI expansion slots. Designed to the Flex ATX specification with a footprint of only 9.0 by 7.5 inches, the Desktop Board fits into imaginative, new system profiles while maintaining backward compatibility with existing MicroATX- and ATX-based chassis designs. Integrators can explore many different design philosophies, including All in One, Slim Tower, Low Profile, and space-saving embedded designs.

The Desktop Board's processor scalability and small size encourages a variety of system solutions. It supports Intel Pentium III processors with 100/133-MHz system bus and Intel Celeron processors with 66/100-MHz system bus to create a wide range of price-performance options. Its Instantly Available PC (Suspend-to-RAM) feature provides low-power consumption and constant computing, enabling system integrators to design and ship products that meet and beat global regulatory requirements for energy savings. With the optional Intel® 82562ET 10/100 LAN on board, the Desktop Board creates a cost-effective integrated networking solution for both PCs and industrial designs.

Solid Quality and Leading Technology

Built with genuine Intel components, the D810E2CB Desktop Board takes full advantage of Intel's technological leadership in system architecture to deliver rock-solid quality in design and operation. All boards are tested for efficient and accurate operation, including environmental stress testing that includes extreme temperatures and shock.

The new Intel 810E2 chipset combination allows the Desktop Board to integrate PC essentials on board. It supports up to 512 MB of PC100 compliant memory, uses ATA 100 technology for faster hard disk access, and provides four USB ports for Plug and Play peripherals. The Desktop Board also offers on board direct AGP Intel® video, integrated AC'97 audio with the Crystal 4201 codec, two PCI slots for system expansion, the Intel® Rapid BIOS Boot to reduce boot time, and optional hardware sensors that monitor power supply voltages and thermal conditions to ensure long-term system reliability.

Snap Integration

In addition to offering the best value design, the D810E2CB Desktop Board saves integrators development time from square one by providing all the essentials of PC operation; sound, video, I/O and more, already integrated on board. To further speed time-to-market, the Desktop Board offers an optional Lite Augmentation Kit that includes ATA 100 cables, I/O shields and product documentation specially designed to facilitate ease of integration. Intel ensures the D810E2CB Desktop Board is fully compatible with the world's most popular software applications. All BIOS and driver updates are easily available off the Intel Web site, greatly simplifying integration and after sale support.

Summary

The Intel D810E2CB Desktop Board is a new, true value platform for the creation of innovative PC designs and industrial systems. It provides a small Flex ATX form factor, Pentium III and Celeron scalability, system efficiencies to trim costs, all the features that customers require, and reliable Intel technology and support. The Desktop Board frees integrators to develop a new generation of capable, efficient value systems. The only constraint left is the imagination.

More Info

For more information on the Intel D810E2CB Desktop Board, visit the D810E2CB Desktop Board pages on the Intel Developer Web site.

To learn more about the Intel 810E2 chipset, visit the 810E2 page on the Intel Developer Web site.

To discover how to give your next system the advantages of the Intel D810E2CB Desktop Board, contact your local Intel representative.

Author Bio

Sandeep Brahmarouthu began working for Intel in December 1997, while still a student. He authored a troubleshooting document on Intel® PCI chipsets in August 1998. He spent 13 months in Intel Customer Support and Application Design-in Center and currently is a member of the Intel Sales and Marketing Rotation Program. Sandeep holds a B.S.E.E. from the University of Southern California.

Take Your System's Temperature with TTB

Gavin Stanley
Senior Mechanical Engineer
Desktop Platform Enabling
Intel Corporation

Overview

As desktop PCs grow more powerful, they require an ever greater ability to dissipate and expel heat. At the same time, consumers want their systems to emit a lower level of noise. Today's thermal issues require new tools for developing and validating thermal solutions for Intel® processors and their PC enclosures or chassis.

In the past, Intel has used a proprietary thermal test board to validate thermal solutions. Intel is now planning to make the ATX and MicroATX Thermal Test Board (TTB) available to developers outside of Intel. It will be a valuable tool for developing thermal solutions for the next generation of performance ATX and MicroATX systems.

The Thermal Test Board

The TTB is a flexible thermal development tool for system integrators, system and chassis designers, and heat sink developers. Designed to the ATX or MicroATX motherboard specification, the TTB allows you to thermally map your PC and test the effectiveness of your system's thermal solution.

The TTB lets you vary subsystem powers and fan speeds within the system as well as monitor the total power being drawn from your power supply. With the TTB, sensors that you place throughout the PC collect temperature data. User-friendly control software then lets you perform extensive thermal evaluations without costly and complicated instrumentation and data acquisition systems.

Thermal Testing

A difficult problem in today's thermal engineering is acquiring complete and accurate information about the thermal capability of a PC platform. Component heat sinks can easily be tested outside the PC enclosure. However, the overall thermal environment in which the heat sinks reside also needs to be understood. The ultimate goal is to determine the heat sink's ability to meet the component's thermal specification when installed in the PC system.

Temperature and air velocity profiles around components can be significantly affected by the system layout, chassis design, fan selection, and power supply selection. Accurate thermal validation for today's PCs means you must validate the thermal capability of not only your heat sinks, but also of your chassis.

The TTB lets you conduct real-time thermal evaluations of both component heat sinks and the PC chassis. With the TTB, you can run tests manually or as automated sequences of tests (batch mode). For example, you can set the TTB to incrementally increase the power levels to particular test vehicles over a set range of tests to avoid cold starts and help shorten test times.

The TTB automates the data collection process and then generates useful reports. TTB results are, in effect, a thermal map of the given configuration of the system or component under test. All TTB tests and results can be saved, copied, or accessed by a standard PC that is connected to the TTB and is collecting the data from your tests.

Thermal Test Board Hardware

The TTB includes many thermal components and test vehicles to provide power to the sub-systems being evaluated in the PC under test. Besides the software for controlling the TTB, the hardware includes:

- Thermal test board (motherboard)
- Thermal test vehicle for the Intel® Pentium® 4 processor
- Thermal test vehicle for the chipset
- Thermal load card for the AGP slot
- Thermal load card for each PCI slot
- Thermal load card for the processor voltage-regulation area
- Thermal load card for each memory module
- Thermal sensors (which plug into connectors on the motherboard) for taking temperature measurements in the system
- Interface cables that connect the TTB to the data-acquisition PC
- Interface cables for the hard drive, CD-ROM, floppy, etc.

The TTB lets you set various power (or current) levels for each component. For example, you can set a different thermal load across each of the four memory sticks as well as for the processor and chipset thermal test vehicles (TTVs). Other load cards, such as those for the PCI and AGP cards, accept variable settings between 3 and 25 Watts.

The TTB's flexibility lets you test to a minimal power configuration, a standard configuration, or even to extremes. For example, to test a system for an AGP card drawing maximum power, you could set the load on the AGP card to 25 Watts. You can then easily monitor the resulting thermal environment (from 0 to 100 degrees C) in the system with the TTB temperature sensors.

The processor thermal test vehicle lets you test CPU power levels out to Intel's predicted end-of-life values rather than only to today's production silicon levels. To help you set appropriate load profiles, the TTB includes some guidance in the form of predefined baseline settings. These predefined load profiles give you a baseline for testing that should be similar to standard loads in existing product chassis. The baseline configurations help you set test variables realistically, so that the resulting thermal map is valuable in identifying and solving any thermal problems that may exist.

Processor Heat Sinks

The TTB lets you specifically test and validate the thermal performance of your processor heat sink. The thermal sensors included with the TTB allow you to make temperature measurements on the processor case (the Tcase) and at the inlet to the heat sink. With this information and the known power level of the processor thermal test vehicle, you can easily calculate the thermal performance of the heat sink under test.

Chassis

The TTB's system-level test evaluates the overall thermal environment in the PC. The data collected in this level of test gives you an overall thermal map for the system under test.

System thermal profiles are affected by many factors: power levels, fan selection and speed, the number and layout of components, and power supplies. Variations in any of these factors affect the overall thermal environment by changing:

- The amount (increase or decrease) in total volumetric airflow through the system, which directly affects air temperature
- How air is directed, blocked, and/or flows through the components, which depends on the configuration and population of drives/cables/cards
- Which components get cool or hot air due to subsystem heating

With the TTB you have the ability to manipulate and evaluate the system's thermal profile while varying:

- The power level of the processor, chipset, PCI cards, AGP card, and memory sticks
- The voltage level for each fan
- The total population of fans in the system
- The power supply selected
- The configuration and population of the drives (hard drive, CD-ROM, floppy drive, etc.)

Power Supplies

A secondary problem with today's more powerful PCs is judging the load that the power supply must deliver to the system. The TTB does not include a power supply. Instead, it uses your power supply to provide power for the thermal tests. The TTB is also able to check the output power on each voltage rail of the power supply. This allows you to check your actual power draw against the power specification of your power supply.

Software Interface

Power levels, fan voltages, and test information are presented through the TTB's Web-based software package. This software program resides on a standard PC connected to the TTB. The software provides clear, step-by-step instructions for setting up tests, running tests, and storing results and test information. It also alerts you if you are exceeding self-imposed power limits or thermal values. Detailed help files are also just a click away.

Summary

Intel's TTB is a tool for validating the thermal requirements of next-generation, Pentium 4 processor performance desktop PCs. By controlling the power consumed by the TTB's components, you can more accurately evaluate a system's thermal performance.

The TTB gives you valuable empirical thermal and power data about a system early in the design process. The TTB also reduces or even eliminates the time-consuming and difficult setup of complicated and expensive lab equipment currently used for thermal testing. Ultimately, the TTB will help speed system development and so reduce your product's time-to-market.

More Info

The TTB will be described in detail at the Intel Developer Forum Conference Spring 2001 (February 26 through March 1). Information about the TTB, such as firmware, front-end software, TTB functionality, and the board itself, will be presented in a general session and a two-hour lab. To register for the Conference, visit the IDF Web site.

Processor information can be found at Intel's Developer Site

Platform specifications and thermal information can be found at the Platform Development Support Web site

Author Bio

Gavin Stanley has been with Intel for three years and has worked extensively on the thermal and mechanical enabling of FlexATX, MicroATX, and ATX system platforms. His most recent work focuses on the system level thermal design and analysis of Pentium 4 processor-based performance platforms. Gavin has presented sessions and labs on thermal enabling at four Intel Developer Forum Conferences. He received his M.S. in mechanical engineering from the University of Washington.

Initiatives and Technologies

CNR Audio Solutions Reduce Costs for OEMs

Philip R. Lehwalder
Senior Hardware Design Engineer
OEM Platform Solutions Division
Intel Corporation

Overview

In this era of multimedia and streaming media applications, the PC audio solution is a critical component in creating the end-user's experience. In the past, motherboard audio designs revolved around expensive and complicated Industry Standard Architecture (ISA) or Peripheral Component Interface (PCI) accelerated (DSP-based) solutions, which were costly from both a bill of materials and a motherboard real-estate standpoint.

OEMs now have a method for constructing an audio solution that costs less and is more flexible than using PCI cards. Chipset vendors have begun integrating high attach rate interfaces directly into their chipsets. One of the interfaces typically included is the Audio Codec '97 (AC '97) interface, shown in Figure 1. With the integration of AC '97 into the chipset, the OEM immediately has a large portion of the overall audio solution.

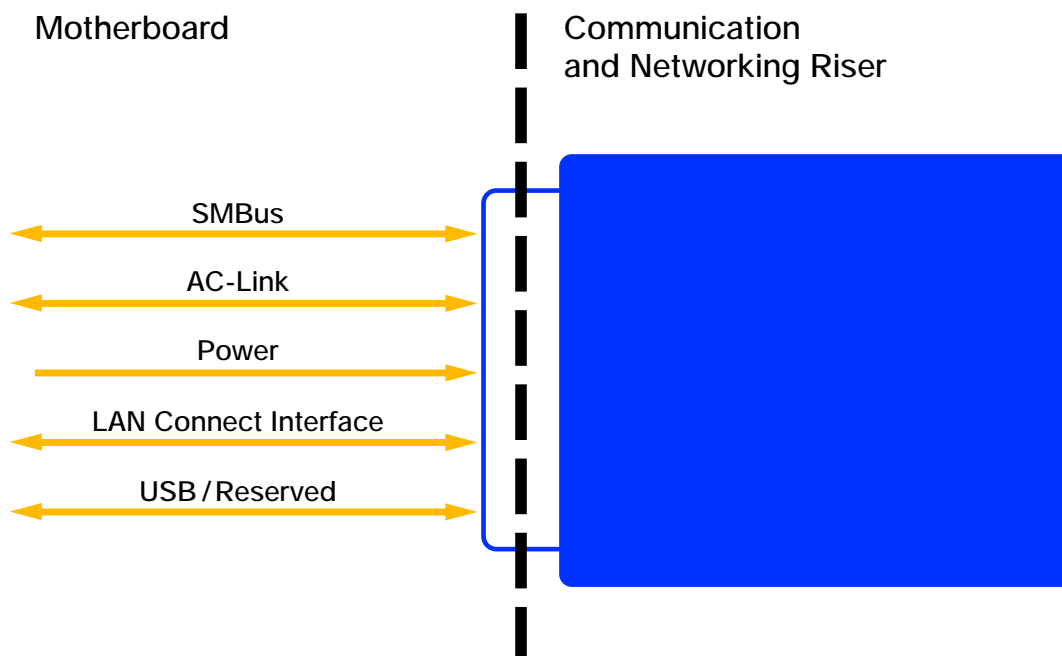


Figure 1. CNR Overview

The other part of the solution comes from using a CNR (Communication and Networking Riser) card that, together with the AC '97 interface and software, provides the same performance as a PCI card. The PC industry has developed CPU-based digital signal processing software that provides feature sets equivalent to popular PCI add-in cards. By using CNR and CPU-based signal processing software, the system integrator or OEM has the ability to supply a scalable audio solution ranging from basic to high-end audio at a price point at or less than that of the PCI add-in card. This allows the OEM or system integrator to maintain or increase sell-up margins relative to PCI add-in audio cards, and continue support of existing business models.

CNR Audio Solutions Provide Value-added Functionality

The CNR specification provides a flexible and cost-effective way to implement an audio solution. The CNR specification is an open industry standard for desktop PC motherboard designs. The idea behind this specification is to provide motherboard designers, system integrators, and OEMs with a robust and low-cost architecture for utilizing the functionality of interfaces that are integrated in the chipset down on the motherboard.

CNR incorporates multiple motherboard interfaces on a single connector, so communication, networking, and audio building blocks can be easily incorporated onto a single card. This provides OEMs, manufacturers, and system integrators with the flexibility to integrate audio functionality as a value-add. It also enables them to develop multi-function cards for specific market segments, such as audio plus 10/100 Ethernet for the corporate market, or audio plus modem for the consumer market.

High-Speed Software Eliminates Need for Accelerators

In the past, PCI hardware accelerators have been used to ensure that PC audio was free of aberrant effects, such as audio dropouts due to variation in system bandwidth caused by other devices within the PC. In addition, many of these PCI accelerators have been capable of performing various effects such as reverb, flanging, chorus, sample rate conversion, and virtual surround technologies such as head-related transfer functions (HRTF), 3D positional, and environmental audio. A typical PCI/AC '97 hardware accelerated audio solution is shown in Figure 2. When the user plays back a WAV file, the data is retrieved by the PCI controller, processed by the accelerator's DSP, converted to AC '97 format, and sent over the AC '97 interface to the CODEC. The audio CODEC then converts the PCM audio data to an analog output, which is passed to the back or front panel jacks, and ultimately to the PC speakers.

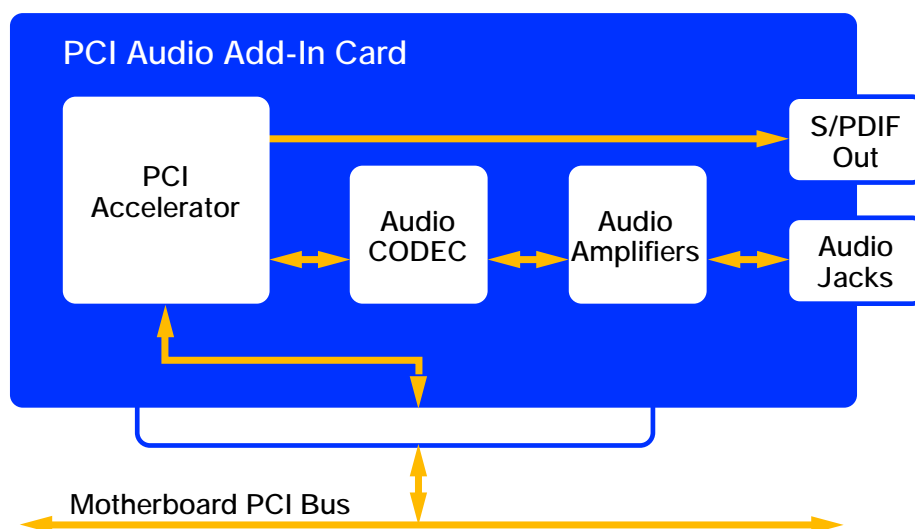


Figure 2. Standard PCI Audio Card Architecture

Today's high-speed processors have the horsepower needed to perform these DSP effects in software without excessive use of processor bandwidth. The development of CPU-based audio digital signal processing software, in combination with CNR, enables a new set of audio hardware configurations.

CNR-based Solutions Eliminate Need for Multiple SKUs

The CNR is a value-add to the OEM or system integrator because it allows multiple audio solutions to be designed for a single motherboard SKU. Previously, the OEM had to manage multiple SKUs of a given motherboard when different audio options were required. OEMs would, for example, take an audio and non-audio SKU for a given motherboard. This allowed the OEM to productize a system with basic audio down on the motherboard and a system with a high end PCI audio solution on a non-audio down SKU of the same motherboard. CNR allows for these options plus several audio upgrade options with a single SKU of the motherboard.

Figures 3a & 3b shows some examples of possible CNR and motherboard/CNR audio solutions. Note that the CNR can also be used as an OEM or system integrator upgrade path for audio where audio down on the motherboard is used in conjunction with a CNR audio solution to create a scalable multi-channel solution that enhances the end-users' experience.

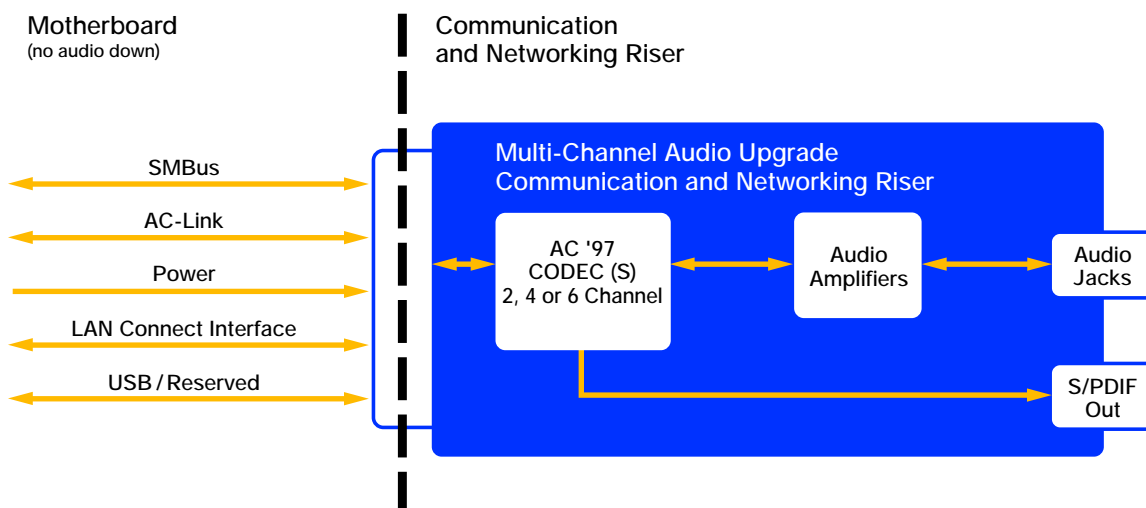


Figure 3a. Audio on CNR and No Audio Down on the Motherboard

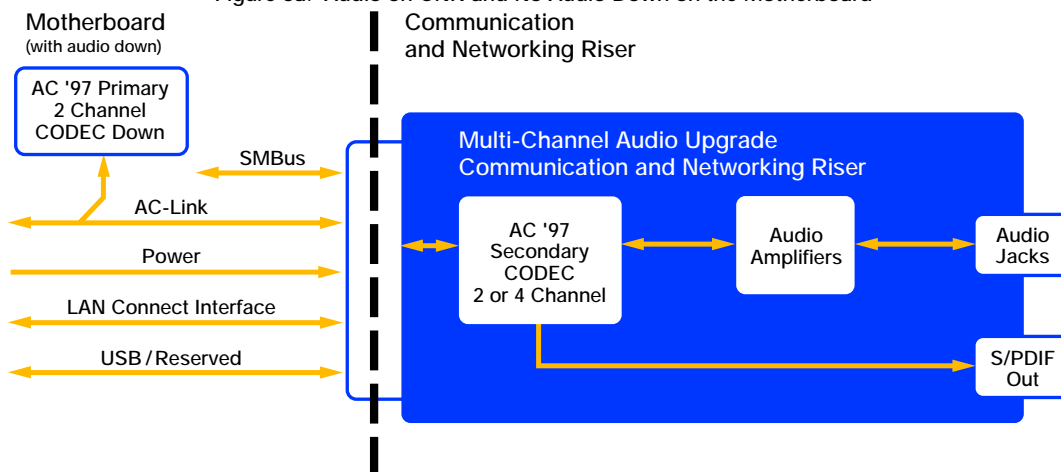


Figure 3b. Multi-Channel Audio Solution Split Between CNR and the Motherboard Audio Solution

Summary

CNR solutions utilize the motherboard AC '97 interface, along with processor horsepower and software-based digital signal processing software, to provide an overall solution that is equivalent to PCI add-in cards at or below the PCI add-in card price point. OEMs benefit by taking full advantage of the AC '97 interface they are already paying for when they purchase a chipset, instead of paying additional cost for PCI add-in audio cards. In addition, the number of motherboard SKUs is reduced, making inventory management much simpler.

More Info

For more information on CNR-based audio solutions, visit:

- Intel's Communications and Networking Riser news Web site
- An overview of CNR technology and CNR's place in the retail market
- Questions and answers about the CNR specification

Author Bio

Philip R. Lehwalder joined Intel in 1995 as a college intern for the Military Special Products Division. He then moved directly to OPSD, where he worked on motherboard audio and modem design. Philip holds an electrical engineering degree from the University of Washington.

IDF Spring Conference Takes Broad Industry View

Vincent Merrick
IDF Conference Content Manager
Intel IDF Group
Intel Corporation

Overview

The Intel Developer Forum (IDF) Conference Spring 2001 is a four-day technical conference that covers a comprehensive range of hardware and software topics for client, server, and communication technologies. This conference offers an opportunity for developers to receive in-depth training in a variety of technologies, and also establish beneficial relationships with developers in other industries. The IDF Conference Spring 2001 will be held in San Jose, California, USA, from February 26 through March 1.

The IDF Conference features several keynotes on the direction of leading-edge Internet technologies. These keynotes are presented by high-profile Intel executives and other influential industry leaders. Intel President and CEO Craig Barrett kicks off the conference by sharing Intel's vision of net architectures for next-generation clients, servers, and communication networks. Other keynote topics include Intel's direction for client and server platforms, advances in wireless and optical networking technologies, and state-of-the-art technologies for enterprise computing platforms. A complete list of currently scheduled keynote topics is located on the IDF Web site.

The IDF Conference Spring 2001 includes 24 tracks of technology demos, overviews, case studies, and in-depth nuts-and-bolts training. Over 200 sessions will cover hardware and software aspects of the latest and next-generation products, tools, and technologies. The conference also includes many labs, which are repeated to help attendees avoid scheduling conflicts. Technical sessions and labs are taught by key Intel engineers and other industry experts. Registration is available online or via phone, fax, or mail, and discounts apply through February 2.

Conference Focus

The IDF Conference Spring 2001 focuses on three main areas: clients, servers, and communications. Sessions in each of these areas answer questions about platform architectures, describe implementations, resolve ambiguities in specifications, explain converging technologies, and offer in-depth training in new tools and technologies.

This conference is not just about Intel, but about industry involvement. With a variety of presenters from Agere Systems (formerly Lucent Technologies) to Unisys, the IDF Conference reaches across many industries boundaries and shows how different technologies can work together. The Telco track, for example, includes a session on InfiniBand* implementation in a digital telephony application. The IDF Conference also offers sessions for complementary technologies, such as Serial-ATA or 1394b bus technologies and test and measurement. Overall, about 40 percent of conference sessions have industry participation.

Although the IDF Conference was originally known for focusing on hardware developments, the conference now offers extensive sessions for software developers. This spring, almost half of all Conference sessions are devoted to software tools for wireless connectivity, optical communications, cellular multiprocessing (CMP), applied computing, network management, code optimization, and other software topics.

Conference sessions also reflect the trend toward open-source code. For example, this conference offers sessions on both proprietary and open-source technologies, from companies such as Microsoft and Red Hat, Inc. For example, IBM will be using a session to explain the Linux* application environment (LAE) and how these open-source applications can run on IBM's proprietary PTX* operating system (OS).

Client Development

Client development tracks offer the most up-to-date information available on high-performance desktop motherboards, platforms, and integration. Some sessions will explore emerging technologies and standards for next-generation platforms and mobile systems. Other sessions and case studies cover software and Web development, security and trusted platforms, applied computing, and ease of use. The Conference will tell how new design practices and technologies developers take ease of use to new heights with wireless computing, home networking, and power management (IAPC), for example.

Server Development

Conference sessions will tie advances in IA-32 and IA-64 technologies to new architectures, such as cellular multiprocessing (CMP) for large-scale environments. One interesting note about this track is that it is completely owned and driven by a third party—once again demonstrating the high level of industry involvement in this conference.

Other server tracks include InfiniBand architecture and networking, enterprise core technology, and topics such as the extensible firmware interface (EFI). The EFI sessions explain how to decouple the hardware platform from the OS and solve traditional BIOS boot issues. Board, system, and software designers will find sessions and case studies on power and chassis elements, memory subsystems, thermal solutions, server management hardware and software, network attached storage, Gigabit Ethernet on motherboards, and many other topics. Other sessions will cover concept technologies under development for high-density platforms.

Communications

The revolution from circuit switching to packet switching, and other advances are pushing optical networking, wireless clients, and wireless connectivity to new levels. These are the technologies driving third-generation (3G) high-power, highly capable phones that can handle voice and data communications on a single platform. Conference sessions cover evolving wireless standards and technologies and infrastructures and architectures, as well as emerging wireless multimedia technologies. For example, the Intel® Internet Exchange Architecture (IXA) sessions teach you how to use new development platforms to improve time-to-market.

The optical networking track is new for the IDF Conference Spring 2001. In this track, you'll find sessions on carrier-core solutions, forward error correction, and high-speed designs for OC-192. Sessions on wireless clients will provide a hardware and software perspective of the capabilities of new, high-bandwidth devices, from phones to personal data assistants (PDAs).

The Intel® Personal Client Architecture (PCA) is the heart and soul of Intel's wireless client. At the last IDF Conference Intel demonstrated the first processor designed specifically for wireless devices. At the IDF Conference Spring 2001, Intel will demonstrate actual implementations based on its XScale™ microarchitecture. These sessions on wireless clients will cover Intel® PCA, flash memory, wireless multimedia, microsignal architecture, and more—all of which are based on XScale microarchitecture. The wide-reaching implementations already available on this microarchitecture prove just how serious Intel is about entering and leading this industry.

Wireless Standards

Currently, the wireless industry is working with several competing standards: HomeRF*, 802.11, and Bluetooth*. This means that it is difficult to get the benefits of an economy of scale that a standard should produce. During the conference, Intel will be offering direction for how to use and position these three wireless standards today. Intel will also provide direction to developers to facilitate the convergence of wireless connectivity on 5 GHz.

Green Design

The IDF conference is a major forum for technologies training and product announcements. Because the conference reaches so many developers across so many industries, Intel is offering a special session to promote environmentally responsible, or “green,” design. A panel of experts will explain practical ways to optimize technology such as minimal energy consumption, recyclable batteries, earth-friendly materials, and earth-friendly end-of-life policies.

Sign Up Now

The IDF Conference Spring 2001 will be held February 26 through March 1, at the San Jose Convention Center in San Jose, California, USA. To sign up for the IDF Conference Spring 2001, fill out the registration form at the IDF Web site. The site includes specific information on keynotes, overview sessions, technical tracks, labs, and demos. The site also includes information on the San Jose area, hotel accommodations, and evening events for attendees.

Summary

The IDF Conference has become Intel's premier developer event. With over 200 sessions, this conference offers technology overviews, panel discussions, case studies, labs, and in-depth training on the latest hardware and software aspects of leading-edge clients, servers, and communications. In addition to technical training, the IDF Conference offers developers a place to establish relationships with other developers.

The IDF Conference demonstrates a high level of industry involvement. Session presenters come from a wide array of companies and interest areas. And the Conference is regarded as a major venue for announcements by third parties and other industry members. The Fall 2000 IDF Conference generated over 2,000 articles and broadcasts over an exciting range of technologies from many industry leaders. Expect no less from the IDF Conference Spring 2001.

More Info

Information and registration forms for the IDF Conference Spring 2001 are located on the IDF Web site. The site includes the current listing of conference sessions, keynote presenters, and details on related conference activities.

A personal event planner (PEP) will also be available on the site in early February. When you enter a topic that interests you, the PEP returns a list of sessions relevant to that topic. This feature lets you investigate sessions you might want to attend. You can then select the appropriate sessions and print a personal schedule. You can also use the PEP at the conference to print a fully updated schedule of the sessions you selected.

Author Bio

Vincent Merrick is Intel's IDF Conference Content manager. He joined Intel in 1984 and has worked in a number of technical and marketing positions, including design engineer, test engineer, technical marketing engineer, product marketing engineer, and OEM account manager. Vince also spent two years as international customer marketing manager for the OEM Platform Solutions Division.

—End of Intel Developer Update Magazine Issue 17—